DSTO

Asynchronous Single Platform Sensor Fusion

Mark L. Krieg

DEPARTMENT OF DEFENCE
DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION

ASYNCHRONOUS SINGLE PLATFORM SENSOR FUSION

Mark L Krieg

Microwave Radar Division
Electronics and Surveillance Research Laboratory

DSTO–TN–0084

ABSTRACT

Multi-sensor tracking potentially has many advantages over single sensor tracking. This report evaluates the performance of a multi-sensor tracking algorithm, the asynchronous fused Kalman filter, using both simulated and real data from two dissimilar sensors. The real data was collected using a sensor fusion test-bed consisting of two sensors, a pulse Doppler radar and an optical video tracker. The performance of the algorithm has been evaluated under various conditions including clear sky, clutter, multiple targets and intermittent sensor operation. The effect of sensor fusion on the system's robustness to model mismatch has also been investigated.

Published by

DSTO Electronics and Surveillance Research Laboratory
PO Box 1500
Salisbury, South Australia, Australia 5108

Telephone: (08) 8259 5555
Facsimile: (08) 8259 6567

AR No. AR-010-219
June, 1997

APPROVED FOR PUBLIC RELEASE

ASYNCHRONOUS SINGLE PLATFORM SENSOR FUSION

EXECUTIVE SUMMARY

Recent trends indicate that multi-sensor fusion will become an essential component of military platforms, particularly in the presence of electronic countermeasures. A sound knowledge and understanding of sensor fusion techniques is essential for evaluating future developments and acquisitions.

A multi-sensor tracking filter, the *asynchronous fused Kalman filter* (AFKF), has been developed to analyse the operation and performance of asynchronous multi-sensor trackers. This algorithm uses a variable update rate Kalman filter, with a common target process model to represent the target dynamics, and an augmented measurement model to take account of multiple sensor inputs.

A sensor fusion test-bed has been developed using the generic pulse Doppler radar (GPDR) developed by the Microwave Radar Division of the Defence Science and Technology Organisation, Australia. An optical tracking system has been added to the radar, and a separate computing platform installed to fuse the radar and optical measurements.

The AFKF algorithm has been evaluated using both simulated and real data. The simulated data demonstrated the domination of the optical measurements on the angle trackers. Under ideal conditions, almost identical tracking errors are achieved by a Kalman filter using angle measurements from the optical sensor, and range and Doppler measurements from the radar.

Radar and optical data collected using the sensor fusion test-bed was used for further evaluation of the tracking algorithm. Discrepancies between the filter models and the actual data prevented the AFKF from achieving the results obtained with the simulated data. The data revealed that the optical measurement noise increased at close target ranges. The subsequent model mismatch could be overcome by using multiple model techniques.

The AFKF performed well with intermittent measurements, easily switching between single sensor and dual sensor operation. However its performance in clutter and in the presence of other targets was poor. The optical tracker was easily seduced from the target of interest by other targets and background objects. The domination of the optical measurements on the operation of the AFKF caused the system to follow the seducing target, subsequently losing track. Some form of data association is required to overcome this problem.

The effect of sensor fusion on the sensitivity of a tracking filter to model errors has been investigated. Generally it was found that the addition of another sensor only affected the sensitivity when the noise covariance of its measurement model was less than that of the original sensor. Under these conditions, the additional sensor dominated the tracking filter, and the system became more sensitive to model errors.

THIS PAGE IS INTENTIONALLY BLANK

Authors

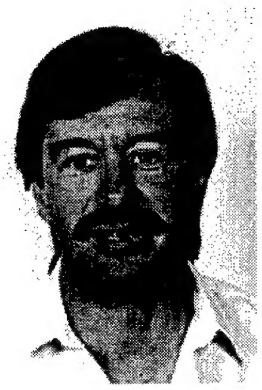

Mark L Krieg
Microwave Radar Division

Mark Krieg graduated with a BE(elect) from the University of Adelaide in 1992, and is currently enrolled as a PhD candidate with the same institution. Mark has worked in various locations within the Defence Science and Technology Organisation, Australia during the past twenty years. Since commencing his professional career in 1992, he has been attached to the Microwave Radar Division, working on the development of an experimental pulse Doppler radar and multi-sensor tracking techniques. His professional interests include radar data processing and sensor fusion.

THIS PAGE IS INTENTIONALLY BLANK

Contents

1 Introduction

Data fusion has generated significant interest in recent times. Although it has only recently gained popularity, its underlying fundamental technologies are well established. It is more accurately described as a new application of existing technologies, rather than a new technology. One particular application of common interest to both civilian and military users is multi-sensor tracking.

A multi-sensor tracking algorithm has been developed to fuse measurements from multiple collocated dissimilar sensors (eg radar, and infrared). This algorithm was designed to operate using measurements of different types, eg angle only, angle and range, etc, from unsynchronised sensors. The measurements arrive from different sensors at different times, at different rates, and possibly intermittently.

A test-bed has been developed for multi-sensor data collection and algorithm testing. It contains two sensors, a pulse Doppler radar and an optical video tracker. Further sensors, such as an ESM receiver, may be added.

A data fusion technology base has been established through collaboration between various divisions of the Defence Science and Technology Organisation, Australia (DSTO), the Cooperative Research Centre for Sensor Signal and Information Processing (CSSIP), and various Australian universities. This technology base will provide the Australian Defence Force (ADF) with local support for the effective evaluation of future developments and acquisitions.

2 Defence Significance

Many ADF platforms already contain multiple sensors with varying degrees of sensor integration and fusion. Others are being upgraded with additional sensors and processing capability. Sensor fusion has the potential to enhance the performance of locating and identifying targets in real, hostile and cluttered environments.

The utilisation of sensors is heavily mission dependent, making the ADF's requirement unique. It may therefore be necessary to develop or modify sensor systems specifically for the Australian environment using local resources, requiring a sound knowledge and understanding of data fusion techniques.

3 Sensor Fusion

3.1 Overview

Data fusion is the combination of information from multiple sources to obtain a more accurate and complete picture of the environment than is possible from a single source. *Sensor fusion* is the combination of data from multiple sensors, and is used in a variety of applications including tracking, identification, command and control, manufacture and assembly, surveillance and robotics. This variety of applications leads to the use of a range

of technologies, including filtering, artificial intelligence, knowledge-based systems, neural networks, estimation theory and wavelet transforms.

Fusing sensors of similar types improves accuracy and reliability through redundancy. Fusing sensors of different types provides additional information, giving a more complete picture of the environment. For example, when fusing radar and optical data, the higher angular resolution of the optical sensor is combined with the range information from the radar. Fusion also provides a mechanism for sensor management tasks, such as sensor queuing and system optimisation. Combining sensor information at a higher level presents an overall picture and assists decision making. In the future this may lead to automated decision and control.

Data fusion can be viewed as a hierarchical structure, as illustrated in the proposal by the Data Fusion Sub-panel (DFS) for the Joint Directors of Laboratories (JDL), Technical Panel for C3 in the United States Department of Defense (figure 1). This proposal introduces three successively higher levels of fusion, with sensor management providing feedback for sensor control.

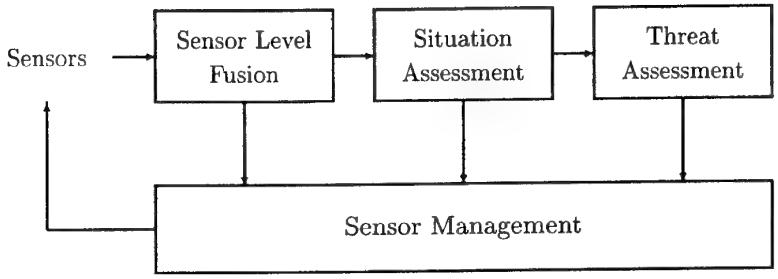

Figure 1: US-JDL data fusion levels.

Sensor level fusion (level 1) refers to directly processing the measurements or information from the sensors. This level includes tasks such as detection, parameter estimation, tracking and identification.

Situation assessment (level 2) processes the outputs from level 1, eg target tracks, and determines the relationship and interaction between individual objects. These relationships and interactions, together with the positional information from level 1, create an overall picture or model of the environment, in particular what the objects are doing and how they are doing it.

Threat assessment (level 3) analyses information from level 2, determining the expected outcomes and possible events arising from the situation. It considers why objects are behaving in a particular way, and what is the significance of this behaviour. This information is used to make decisions and initiate responses.

Sensor management uses the outputs from the above levels to control and configure the individual sensors from a system point of view. This optimisation is dependent on the particular application and operating environment.

This report addresses the level 1 function of target tracking, or more specifically, the fusion of measurements from multiple collocated sensors into a single track. This task, as

with most level 1 fusion tasks, may be separated into three separate sub-tasks, namely positional fusion or registration, data association and data combination (figure 2).

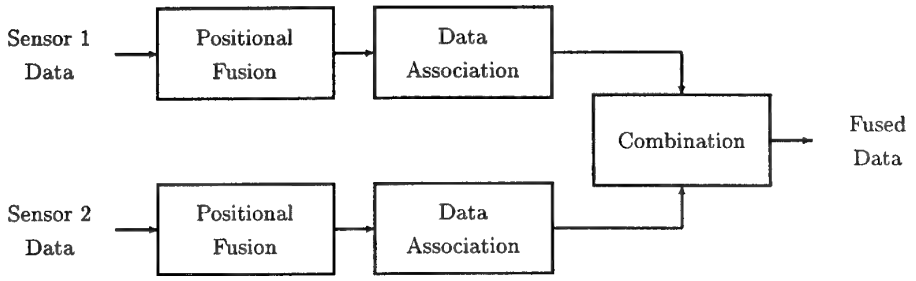

Figure 2: Fusing data from two sensors.

Positional fusion aligns the data both spatially and temporally. This guarantees that all sensors are looking at the same point in space and time. This task is generally simpler for single platform configurations, because the relative positions of all sensors are known, and reliable temporal synchronisation of the sensors is easily implemented.

Data association determines which measurements originate from a target, and therefore may be used for tracking. Nearest neighbour, Probabilistic Data Association (PDAF), Fuzzy Logic and Dempster-Schafer are examples of techniques that have been used in attempts to solve this problem. This is perhaps the most difficult part of the fusion process, and research is far from exhausted.

Data combination is the combination of the data into a single entity. This is often a weighted sum, with the weights based on some optimisation criteria, such as least squares or maximum likelihood. In some cases the combination may be rule based, or simply a case of selecting the data which is best in some sense.

Multi-sensor tracking provides a number of advantages over single sensor tracking. These may include:

1. improved track accuracy through availability of additional and possibly more accurate measurements,

2. greater reliability through sensor redundancy,

3. increased robustness to noise, clutter and countermeasures, and

4. stealth operation through reduced active emissions (ie active sensors only transmit when absolutely necessary).

However these advantages must be balanced against greater system complexity and cost.

4 Previous Work

The majority of work on multi-sensor tracking has centred on fusing simultaneous measurements from two or more sensors. This work has been well documented in the open literature, with [1, 2, 3] providing useful reading on the subject. Little emphasis has been placed on the specific problem of asynchronous fusion. One exception is the work of Blair, Rice, Alouani and Xia [4, 5], who addressed the problem of asynchronous fusion for target tracking. They extrapolated the measurements to a common time before combining the data. Extrapolating reduces the computational complexity, but also reduces the track accuracy. In many applications, the availability of increased computing power has removed the need for extrapolation.

The problem of sensor registration, often ignored in the past, has been receiving increased attention. In particular, techniques for solving this problem automatically are being pursued at the CSSIP [6].

The use of a Kalman filter with augmented measurements and a common dynamic model, has proven a useful tool for combining measurements from different sensors. However, the most difficult problem is data association. This problem is common to single sensor multi-target tracking, and has attracted considerable interest in the scientific community.

Nearest neighbour, ie using the measurement located closest to the predicted target position, is the simplest form of data association. Alouani, Rice and Auger [7] developed a generalised nearest neighbour algorithm that uses a weighted sum of all measurements within a validation gate. These weights are determined by the distance between the measurement and the predicted target position.

Bar-Shalom and his team at the University of Connecticut have developed the Probabilistic Data Association Filter (PDAF) [2]. This filter probabilistically weights all measurements within a validation gate according to their distance from the predicted target position. The resulting composite measurement is used to correct the tracking filter. The equivalent multi-target tracker is known as the Joint PDAF, but is computationally very expensive. Various sub-optimal alternatives have been developed. PDAF methods are particularly suited to tracking targets in clutter.

Houles and Bar-Shalom [8] developed a multi-sensor tracker to track manoeuvring targets with a radar and infrared sensor. This algorithm combines the PDAF with an interacting multiple model (IMM) tracker. The PDAF handles the uncertainty in measurement origin, and the IMM gives the algorithm its manoeuvre handling capability.

Evans and others [9] at the University of Melbourne have developed new variations on the PDAF. One such variation is the Integrated PDAF [10], which treats track existence as an event with a certain probability.

Deb and others [11, 12] partition the problem into a set of feasible measurement-to-target associations. They choose the most likely partition by maximising the ratio of the likelihood of all measurements, given the partition, to the likelihood of all measurements being false alarms.

Blackman's extensive experience in association and fusion of multiple sensor data includes multi-hypothesis tracking (MHT) [13, 14]. This algorithm considers every possible

measurement to target assignment, and therefore grows exponentially. To overcome this impracticality, least likely hypotheses are removed. Blackman has developed a variation, called most probable hypotheses tracking, where all significant hypotheses are considered.

Avitzour [15] used a maximum likelihood technique to solve the data association problem. He treats the problem as one of missing data, where the measurement to target associations are the missing data. Streit and Luginbuhl [16] have combined probabilistic and maximum likelihood techniques in their Probabilistic Multi-hypothesis Tracking Algorithm (PMHT). This algorithm operates in a single sensor multi-target environment, estimating the assignment probabilities and the target states simultaneously. This work has been extended to multi-sensor tracking [17, 18, 19].

Hong [20, 21] has worked on adaptive sensor fusion. He assumes unknown noise covariances, and uses the tracking filter innovations to determine the optimal gains.

Xie and Evans [22] have used hidden Markov model tracking for multiple targets. They use the concept of a mixed track. The individual tracks are easily identified in the mixed track. Martinerie and Forster [23, 24] have developed a two stage fusion algorithm. The first stage determines a distribution of likelihoods over a discrete space. The second stage uses hidden Markov model techniques to discern the most likely track.

5 Algorithm Development

5.1 Requirement

In a real sensor suite containing multiple sensors of different types (eg radar and infrared), it is unlikely that the measurements from all sensors will arrive at exactly the same instant in time. Different sensors are likely to have different and unsynchronised measurement update (sample) rates. Non-uniform rates may arise from missed measurements, caused by either faint targets or deliberate emission control. Dissimilar sensors may also produce different types of measurements, eg bearing and range from a radar, and just bearing from an infrared device. I will refer to sensor fusion under these conditions as *asynchronous sensor fusion*, or more simply, as *asynchronous fusion*.

The sensors are assumed to be collocated on the same mount. As the spatial relationship between the sensors is fixed and known, the sensor registration problem (positional fusion) is trivial and will not be addressed in this report.

5.2 Tracking Algorithm

5.2.1 Kalman Filter

One of the most popular tracking filters is the *Kalman filter* [1]. It uses two linear-Gaussian models, a process or dynamic model to represent the target dynamics, and a measurement model to represent the measurements received from the sensor. The dynamic model describes how the target state changes between measurements, and is expressed as

$$x_{t_{i+1}} = F_{t_i} x_{t_i} + w_{t_i} \tag{1}$$

where at time t_i, x_{t_i} is the target state, F_{t_i} is the transition matrix describing the target dynamics from time t_i to t_{i+1}, and w_{t_i} is a zero mean uncorrelated Gaussian process noise with covariance Q_{t_i}. The process noise models deviations of the target from the expected dynamics, eg manoeuvres. t_i denotes a particular time instant, where $t_i \geq t_j$ for $i > j$.

At each filter update time, t_i, the measurement model is defined as

$$z_{t_i} = H_{t_i} x_{t_i} + v_{t_i} \tag{2}$$

where z_{t_i} is the measurement, H_{t_i} is the measurement matrix, and v_{t_i} is a zero mean uncorrelated Gaussian measurement noise with covariance R_{t_i}. The measurement matrix represents the relationship between the measurement and the target state. The measurement noise includes environmental and sensor noise in addition to any mismatch between the measurement model and the sensor.

The filter produces an estimate of the target state at each measurement or sample time. The state estimate obtained at time t_i using all the measurements up to time t_j is denoted as $\hat{x}_{t_i|t_j}$, and is obtained using a two step process. The first step is the prediction step, where the previous state estimate and the dynamic model are used to predict the new state of the target, $\hat{x}_{t_{i+1}|t_i}$. This estimate and its error covariance, $P_{t_{i+1}|t_i}$, are determined using the following equations.

$$\begin{aligned} \hat{x}_{t_{i+1}|t} &= F_{t_i} \hat{x}_{t|t} \\ P_{t_{i+1}|t} &= F_{t_i} P_{t|t} F_{t_i}^T + Q_{t_i} \end{aligned} \tag{3}$$

The second step is the correction step, where the measurement received at time t_{i+1} is used to correct the predicted state estimate. This new estimate, $\hat{x}_{t_{i+1}|t_{i+1}}$, and covariance, $P_{t_{i+1}|t_{i+1}}$, are determined as follows.

$$\begin{aligned} K_{t_{i+1}} &= P_{t_{i+1}|t_i} H_{t_{i+1}}^T \left(H_{t_{i+1}} P_{t_{i+1}|t} H_{t_{i+1}}^T + R_{t_{i+1}} \right)^{-1} \\ \hat{x}_{t_{i+1}|t_{i+1}} &= \hat{x}_{t_{i+1}|t_i} + K_{t_{i+1}} \left(z_{t_{i+1}} - H_{t_{i+1}} \hat{x}_{t_{i+1}|t_i} \right) \\ P_{t_{i+1}|t_{i+1}} &= \left(1 - K_{t_{i+1}} H_{t_{i+1}} \right) P_{t_{i+1}|t_i} \end{aligned} \tag{4}$$

In many applications measurements arrive at equally spaced time intervals, and the Kalman filter is implemented with a fixed update rate. When its update interval, or time between measurements, is not constant, it is known as a *variable update rate Kalman filter*, and in this form is particularly suited to asynchronous sensor fusion. An in depth discussion of the derivation of the Kalman filter for both fixed and variable update rates can be found in [25].

5.2.2 Asynchronous Fused Kalman Filter

The *Asynchronous Fused Kalman Filter* (AFKF) is a multi-sensor variable update rate Kalman filter with a single state space dynamic target model, and a time variant augmented measurement equation. At any measurement time, the augmented equation consists of a stack containing the individual measurement equations for each sensor providing measurements at that time.

The radar and optical systems measure target position in azimuth, elevation and range, ie radar coordinates. The resulting dynamic model is non-linear, with complex dependencies between these positional coordinates and their derivatives. For non-manoeuvring targets or targets at long range, the non-linearities and inter-relationships can be modelled as additional process noise. Therefore the dynamic model becomes linear, with a state vector denoted $(\eta,\ \dot{\eta},\ \ddot{\eta},\ \varepsilon,\ \dot{\varepsilon},\ \ddot{\varepsilon},\ R,\ \dot{R},\ \ddot{R},)^T$, where η represents azimuth or bearing, ε represents elevation, and R represents range. ($\dot{a}$, and $\ddot{a}$ indicate the first and second derivative of a respectively.)

The AFKF uses measurements from either the radar, optical sensor or, if the measurements happen to arrive simultaneously, both together. It does this by allowing the measurement model (2) to vary over time. If the measurements arrive from the radar, the values of H_{t_i} and R_{t_i} represent the measurement function and noise covariance of the radar, and likewise if the measurements originate from the optical sensor.

If measurements arrive from both sensors simultaneously, the measurements and their models can be stacked, ie

$$z_{t_i} \equiv \begin{bmatrix} z_{t_i}^{(r)} \\ z_{t_i}^{(o)} \end{bmatrix} = \begin{bmatrix} H_{t_i}^{(r)} \\ H_{t_i}^{(o)} \end{bmatrix} x_{t_i} + \begin{bmatrix} v_{t_i}^{(r)} \\ v_{t_i}^{(o)} \end{bmatrix}$$

$$R_{t_i} = \begin{bmatrix} R_{t_i}^{(r)} & 0 \\ 0 & R_{t_i}^{(o)} \end{bmatrix},$$

where the superscripts (r) and (o) denote the appropriate radar and optical variables respectively. If simultaneous measurements are rare, as is the case for the radar and optical tracker, it is simpler to process them using a zero update interval in the Kalman filter.

5.3 Adaptive Tracking

The Kalman filter requires the covariance of the measurement noise for each sensor, and the covariance of the process noise. Generally the values of these parameters are assumed to be known, and they must match the data for optimal tracking. The sensor noise covariance depends on the sensor characteristics and the environmental conditions. The process noise covariance is determined by the mismatch between the actual target dynamics and the assumed model. As the actual covariances all vary over time and the selected covariances are constant, matched conditions cannot be maintained. (The effect of model mismatch is discussed in section 8.2.)

An alternative is to estimate these parameters from the data. This is known as adaptive filtering, and a number of techniques have been investigated. The simplest, explained by Gelb [26] and Mehra [27], estimates the actual measurement noise and state error covariance from the autocorrelation functions of the measurement innovations (the errors between the measurements and the predicted target state). This method is invoked after filter initialisation, ie when the Kalman filter has reached its steady state condition. Mehra outlines this technique in detail for a constant update rate Kalman filter.

This technique becomes significantly more complex when applied to asynchronous sensor fusion. The system becomes time varying, requiring the storage of previous filter parameters such as gains, covariances, etc, and the autocorrelation estimates for sensors with low measurement rates may be poor.

In practice, not all elements of the state can be directly measured. This coupled with small measurement update intervals (order of tens of milliseconds) may cause over estimation of the process noise covariance. This introduces additional noise into the output track and therefore reduces tracking performance.

Adaptive systems are inherently less stable than fixed systems, increasing the likelihood of the filter becoming unstable and losing track. The above mentioned over estimation of the process noise covariance is a likely cause of filter instability in applications with high measurement rates. Therefore the process noise covariance is allocated an assumed value chosen from the expected target dynamics. During steady steady state operation, the measurement noise covariance is usually much greater than the state estimate covariance. Therefore the measurement noise covariance can be approximated by the innovation covariance. This provides reasonably accurate estimates, and hence good filter stability over a wide range of operating conditions.

6 Testbed

6.1 Overview

A single platform sensor fusion test-bed was developed to collect multi-sensor data, and to test sensor fusion algorithms in real time with real targets. It contains two sensors, a pulse Doppler radar and an optical tracker, both physically attached to the same mount. Additional sensors, such as infrared and ESM, may be added at a later date.

A functional diagram of the testbed is shown in figure 3. R, η and ε denote the range, azimuth and elevation respectively, and errors are indicated by a preceding Δ. A dot above the parameter symbol indicates a rate or derivative, and estimated parameters are denoted by a caret.

6.2 Hardware

The generic pulse Doppler radar (GPDR) is an experimental pulse Doppler radar, operating in the 9–10 GHz region of the electro-magnetic spectrum (see specification, Appendix A). It was developed by the Microwave Radar Division of the Defence Science and Technology Organisation, Australia for radar evaluation and development. The radar is housed in a self contained trailer, and is capable of operating from a portable 3-phase generator (figure 4). The trailer is large enough to house the optical signal processing equipment and the fusion processor.

The radar is almost completely software controlled, allowing most radar parameters and the operator interface to be configured by software. All the radar parameters can be recorded using an inbuilt data logging facility.

A colour CCD video camera, mounted directly above the radar antenna (see specification, Appendix B), is used as the optical sensor. It is housed in a weather proof housing and utilises an automatic iris control to prevent damage by the sun. It is fitted with a 75 mm lens through a focal length doubler, giving an effective focal length of 150 mm. This gives a field of view of approximately 3.4° in azimuth by 1.8° in elevation. This is comparable with the radar's 3 dB beam-with of approximately 1.6°. The two sensors are aligned to maximise their overlap.

The video output from the camera is fed to an automatic video tracker (AVT) housed inside the GPDR trailer. The ADEPT 20 Automatic Video Tracker (see specification, Appendix C) was chosen for its flexibility and level of control. It has been specially modified to provide a confidence measure for its position outputs. The video tracker extracts the target from the video image, and outputs its position relative to boresight. These measurements are available on a frame by frame basis. Access to these measurements, and control of the tracker, is achieved through either a serial link or a VME interface.

A personal computer (PC) is used to control the AVT and display its output. The PC is connected to the AVT through a PC/VME interface obtained from the *Bit3 Corporation*. The VME interface provides greater access to, and control of, the AVT than is possible through the serial port.

A separate PC is dedicated as the *fusion centre*. In addition to performing the fusion function, the fusion centre provides an operator interface, and controls the GPDR antenna mount and range and velocity gates. This control is achieved by inserting the fusion centre into the GPDR's control loop. A 16 bit parallel data bus provides the connection between the sensors and the fusion centre.

Coarse alignment of the camera to the radar antenna is performed as part of the radar calibration procedure. Fine alignment is achieved through the operator interface.

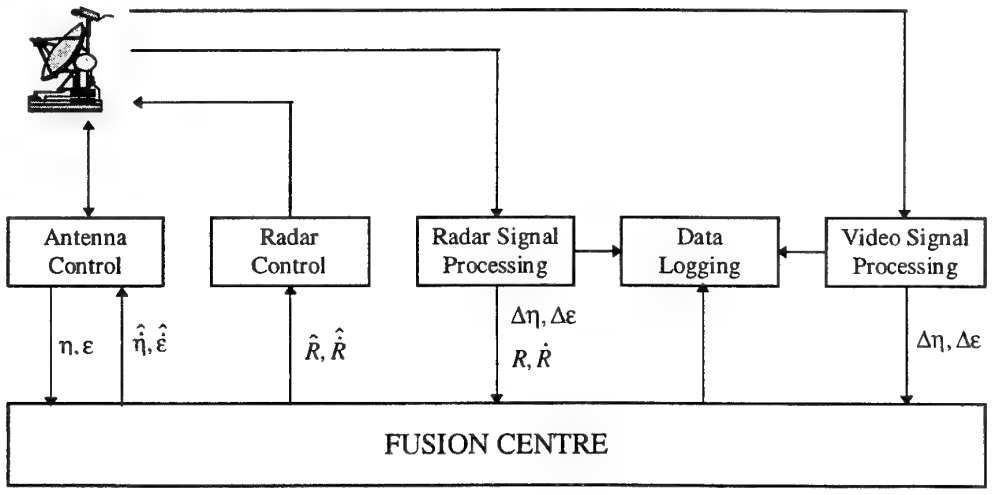

Figure 3: Testbed block diagram.

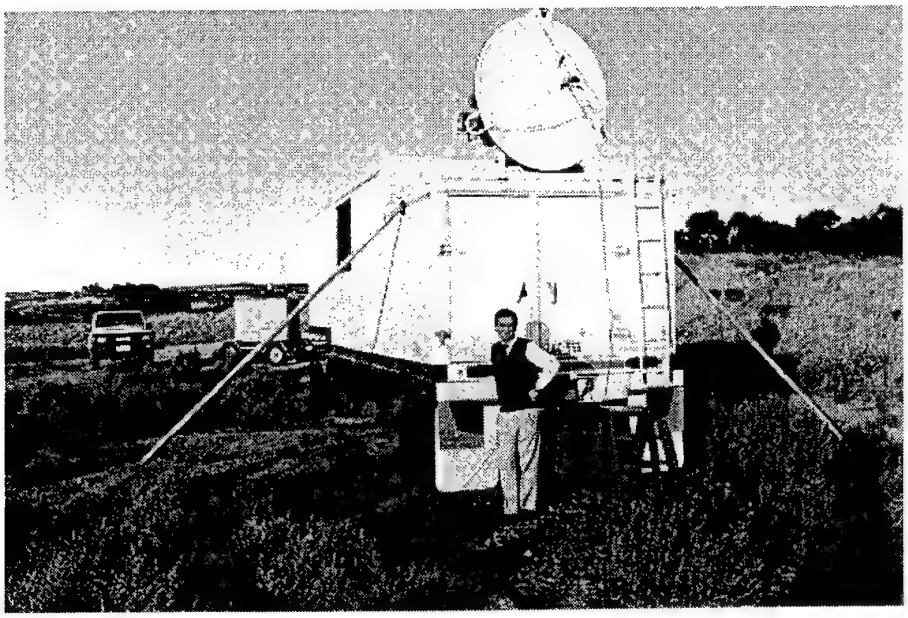

Figure 4: Sensor fusion testbed on location.

6.3 Software

A software library has been developed for the fusion centre. This library provides an interactive environment that allows the operator to control the system and view its status. In particular it allows the operator to:

1. create a new or load a previously saved system configuration;
2. add new sensors to a system configuration;
3. modify the parameters of existing sensors (to match the hardware);
4. connect or disconnect individual sensors;
5. initiate and terminate tracking; and
6. record system and tracking parameters.

The operator display consists of a pull-down menu bar across the top, and a status bar displaying relevant help information along the bottom. The remainder of the screen is split horizontally, the upper section reserved for sensor windows, and the lower containing a parameter display window. Each sensor in the system is identified by a sensor window in the upper display region. This window contains the name of the sensor, its configuration and its current connection status. The lower parameter display window contains the status of various system and track parameters, allowing the operator to observe the system's operation. The operator controls the system using keyboard commands and pull-down menus.

The library also performs a number of background tasks such as:

1. transferring measurements from the sensors to internal queues ready for processing;
2. pulling measurements from the queues in time sequence for processing;
3. transmitting control information to the sensors; and
4. extrapolating the tracking filter in the absence of measurements.

The library does not include the fusion algorithm. This is the responsibility of the system developer, which is incorporated into the system by overloading (or overwriting) various generic functions in the library. Obviously one such function must perform the actual fusion or tracking, but other functions are available for:

1. initialising track and fusion parameters;
2. terminating tracking;
3. drawing the display window;
4. placing parameters into and drawing the display window;
5. sending control parameters to the sensors;
6. selecting parameters to be recorded;
7. selecting a header for the file containing the recorded data; and
8. specifying a system name or identifier.

Other functions, such as those for pulling sensor measurements from the queues and reading the status of system parameters, are available for use by the system developer.

The GPDR software has been modified to include a software interface library. The radar has been configured to operate in either stand alone or fusion mode, the mode of operation being selected by the operator. When used in fusion mode, the radar sends its measurements to the fusion centre. Its servo drives and range and velocity gates are controlled by data received from the fusion centre through the interface bus.

An operator interface to the AVT has been developed. It displays the current tracker outputs and status, and provides a means for the operator to control and configure the AVT through a series of menus and keyboard commands. The operator display is similar to that described for the fusion centre but without the sensor windows. The outputs from the video tracker and its status are displayed in real time. The operator has control over a number of AVT functions including:

1. automatic or manual tracking and detection;
2. choice of tracking algorithm;
3. choice of preprocessing algorithm;

4. tracking window position, size and mode of operation;
5. conditions that must be met before a detection is recognised (eg target speed, direction, size, etc);
6. break-lock and coast operation (in the event of track loss); and
7. tracking filter parameters for filtered outputs if used.

This interface also transmits the measurements received from the AVT to the fusion centre.

7 Algorithm Evaluation

7.1 Overview

The AFKF algorithm has been evaluated using both simulated and real data. The simulated data allowed the algorithm to be tested under tightly controlled conditions, and provided ground truth or true target position. Each track consisted of measurements from two simulated sensors. The real data was collected from various light aircraft and commercial carriers, and contains both radar and optical measurements [28].

The performance of the AFKF was compared to the performance of two single sensor Kalman filters, each using only the measurements from one sensor. Performance comparisons were based on the covariance of the track errors, obtained from the target state estimates and the ground truth.

7.2 Simulated Data

The results from the simulations verified that improved tracking performance, ie lower track error covariance, is obtained by fusing measurements from an additional sensor. The greatest improvement occurs when the measurement noise covariance of both sensors are equal, and the measurements arrive simultaneously from both sensors. When the sensor covariances are unequal, the tracking performance approaches that of single sensor tracking using the sensor with the least noise. In the case of the radar and optical sensor, the optical sensor has a lower noise covariance than the radar. The angular tracking performance of this system is therefore only marginally better than that of the optical system alone. This suggests that only the range and Doppler velocity measurements are required from the radar.

In practice, the radar has a lower measurement update rate than the optical tracker. Simulations under this condition showed a decrease in performance of the radar only Kalman filter. The AFKF's observed performance did not alter because of the much lower noise covariance of the optical sensor. Theoretically, the AFKF's performance would have deteriorated marginally.

The AFKF's performance with manoeuvring targets was investigated. Large errors were encountered during the manoeuvres, caused by significant mismatch between the

data and the algorithm's dynamic model during the manoeuvre. Increasing the process noise covariance of the AFKF reduced these errors, but also decreased the straight line tracking performance. Large errors may result in track loss in real systems, so the tracker design is a compromise between straight line and manoeuvring targets.

The adaptive algorithm (section 5.3) displayed similar results to those obtained from the AFKF. The major difference was an increased track error when targets manoeuvred, caused by an overestimation of the measurement noise covariance during the manoeuvre.

7.3 Real Data

7.3.1 Data Collection

The data used in this study was obtained by locating the test-bed at an elevated site approximately 3km from, and 140 metres above, a light aircraft airfield situated in the northern suburbs of Adelaide, South Australia. This site also provided a clear view of the airspace above the Adelaide International Airport, some 25 km away.

The collected data was arranged into data sets containing radar and optical measurements corresponding to single target tracks. These data sets include aircraft crossing, aircraft landing and taking off, and background optical and radar clutter (buildings, roads, etc). The video tracker was not available during data collection, so the optical measurements were obtained by processing the recorded imagery through the video tracker.

The true target position is required to determine the tracking errors. This ground truth was not available, so the target position was estimated using polynomial fitting techniques on the data. This smoothes the measurements and will remove any short time constant variations, such as those produced by air turbulence. This will artificially increase the measured track error covariance.

During data collection, the absence of the video tracker meant that the sensor mount was controlled solely by the radar. This provided the opportunity for the optical system to re-acquire lost tracks. However if the radar lost track, the optical sensor also lost track.

7.3.2 Results

The AFKF was first evaluated under near ideal, or clear sky, conditions in which clutter was insignificant, and no other targets were present.

The AFKF and both single sensor Kalman filters, ie the radar Kalman filter and the optical Kalman filter, successfully maintained track on the target (aircraft) of interest. The state estimate error covariance, ie *track error covariance*, was used as a measure of tracking performance. This was estimated by subtracting the estimated ground truth from the measured covariance of the state estimate error.

Tables 1 and 2 show typical values of measured track error covariance obtained from the AFKF, and the two single sensor Kalman filters, for both light and commercial aircraft. The measured results do not achieve the expected theoretical values [29], probably due to

Table 1: Track error covariance of a light aircraft.

Tracker	Azimuth (mrad^2)	Elevation (mrad^2)
Radar	0.083	0.18
Optical	0.037	0.083
Fused	0.037	0.072

Table 2: Track error covariance of a commercial aircraft.

Tracker	Azimuth (mrad^2)	Elevation (mrad^2)
Radar	0.082	0.050
Optical	0.053	0.019
Fused	0.050	0.017

a combination of mismatch between the data and the tracking filter models, and errors in the estimated ground truth.

In the case of the light aircraft, the azimuth track errors from the optical Kalman filter and the AFKF are consistent with the dimensions of the target. The greater than expected errors result from the optical tracker algorithm's inability to locate a consistent centroid position on the aircraft body. This problem is particularly evident at close range, where the aircraft occupies a significant number of pixels within the image. This additional measurement noise was not included in the filter models, causing a severe mismatch between the data and the filter models. This mismatch degrades the performance of the optical Kalman filter and the AFKF to a point where their performance is inferior to that of the radar Kalman filter (Table 3). This is caused by the AFKF virtually ignoring the radar measurements, and relying upon the now much noisier optical measurements.

Table 3: Track error covariance of light aircraft at close range.

Tracker	Azimuth (mrad^2)	Elevation (mrad^2)
Radar	0.10	0.021
Optical	0.34	0.044
Fused	0.31	0.040

The commercial aircraft provides an excellent example of one sensor losing the target. In this case, the target moves out of the optical sensor's range of operation, and the AFKF automatically reverts to single sensor tracking using the radar measurements. This is illustrated by the increase in azimuth track error shown in figure 5, where the optical measurements ceased after 125 seconds.

The elevated site provided a good opportunity to evaluate tracking in clutter, as the sensors often operated against an urban background. The optical tracker had difficulty under these conditions, frequently acquiring various objects in the background. To a lesser extent, the radar suffered a similar fate in the presence of large reflectors, eg industrial sheds.

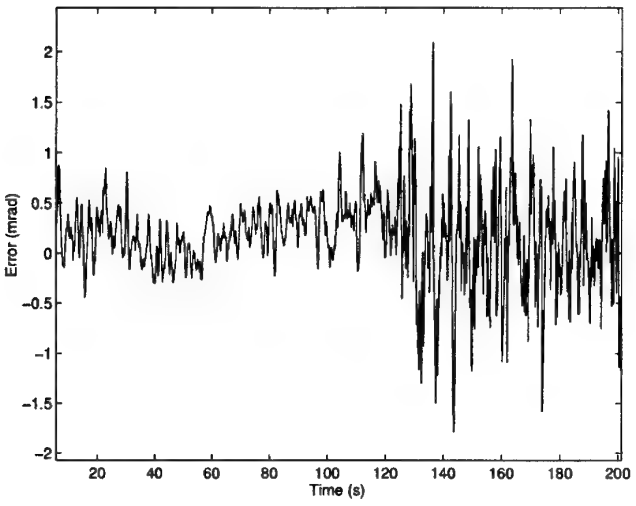

Figure 5: Azimuth tracking error increase due to loss of optical measurements.

Figure 6 shows the azimuth tracking errors obtained from a target which was tracked in clutter. The dotted line represents the tracking errors obtained from the optical Kalman filter, the dashed line the radar Kalman filter, and the solid line the AFKF.

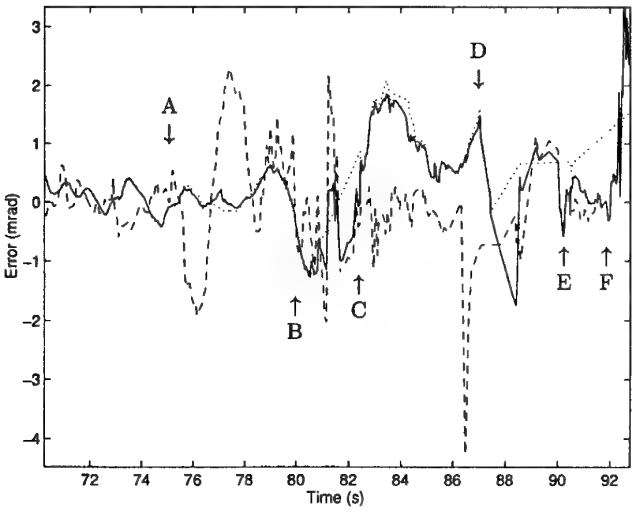

Figure 6: Azimuth tracking errors in clutter.

Initially, all trackers were following the target in a relatively clutter free environment. At 'A' radar clutter was introduced, severely degrading the performance of the radar Kalman filter. The AFKF was not significantly affected because of the dominance of the optical measurements. At 'B' the optical sensor loses the target, and the AFKF only receives radar clutter measurements. At this point the target has been lost. Optical clutter seduces the AFKF between 'C' and 'D'. After a few spurious false alarms from

both sensors, radar measurements recommence at 'E'. However these are not from the target, and cease at 'F'.

A second example is illustrated by the azimuth and elevation tracks in figure 7. In this scenario, radar clutter first appears at 'A'. As in the previous example, the AFKF maintains track using the optical measurements. At 'B' the radar is seduced by clutter, believed to originate from a large iron building near the bottom edge of the main beam. The AFKF continues tracking the target until 'C', when all measurements cease because the mount has moved the sensors away from the target. At 'D' a few spurious radar measurements influence the AFKF and radar Kalman filter until, at 'E', the target fortuitously re-enters the field of view, and is re-acquired by the optical tracker. The AFKF and optical Kalman filter then continue to track the target.

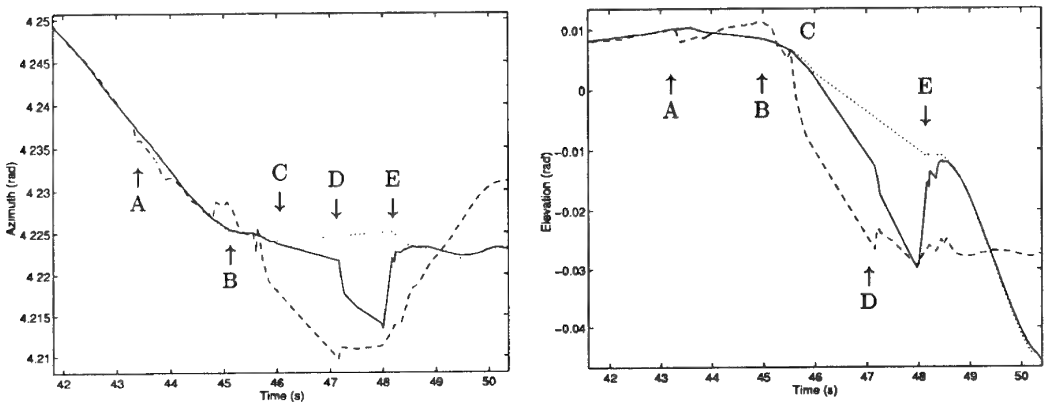

(a) Azimuth.

(b) Elevation.

Figure 7: Tracking in clutter.

In general, it was more common for the optical tracker to be effected by clutter and subsequently cause system track loss. This is because the optical sensor is two-dimensional, and therefore has no range or depth information. Therefore it receives measurements from both the target and clutter. As a result, the AFKF performed poorly in clutter because of the strong influence of the optical measurements on its operation. When optical measurements were available, radar clutter had little effect on the AFKF's performance because the radar measurements contributed little to the fused track.

Other targets passing in front of the target of interest, and large targets in general, often seduced the optical tracker and the AFKF from the target of interest. An example demonstrating this effect is illustrated by the measurements and fused track (dashed line) shown in figure 8. The track commences at 'A' and proceeds in a clockwise direction. The AFKF is seduced from the target by clutter at 'B', and other targets at 'C' and 'D'. These other targets are lost when they leave the optical field of view. After losing the target at 'D', the optical tracker is seduced by several interferences near the top of the field of view ('E') before re-acquiring the target. Normal dual sensor tracking continues until the radar loses the target ('F'), at which point the AFKF tracks the target using only the optical

measurements until the radar re-acquires the target ('G').

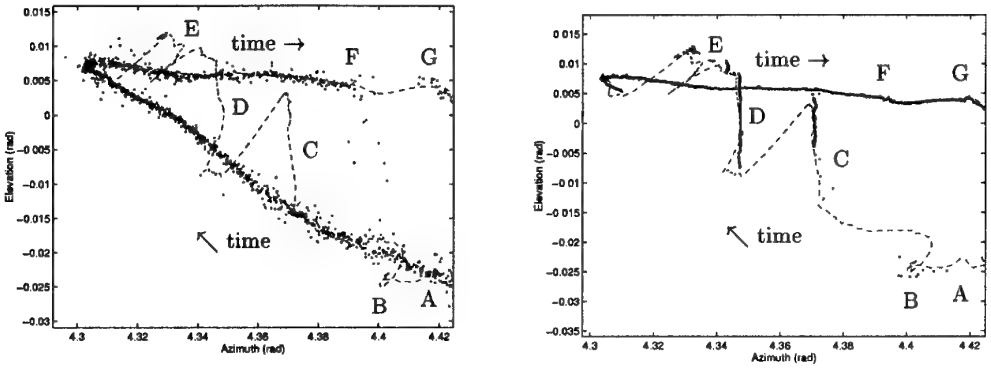

(a) With radar measurements.

(b) With optical measurements.

Figure 8: Azimuth-elevation track with interfering targets.

The radar only (dotted line), and fused (solid line) plots from a second example are shown in figure 9. Commencing from the right, both sensors track the target until the optical sensor is seduced by another optical target at 'A'. The AFKF track follows this target until it is lost by the optical sensor. The AFKF track attempts to re-align itself with the radar measurements, but is prevented by several short term optical interferences at 'B', 'C' and 'D'. Although they only exist for a short period, the dominance of the optical measurements is sufficient for them to have a significant effect on the AFKF track. Therefore once a track is seduced, it may take only relatively minor interferences to prevent re-establishment of the correct track.

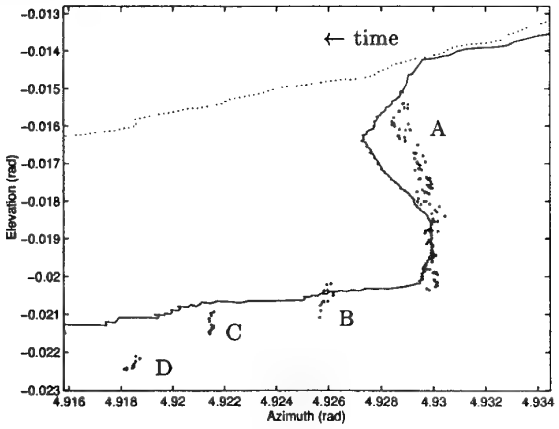

Figure 9: Azimuth-elevation track with interfering targets.

In real time operation, track loss will occur once the optical tracker has been seduced.

This will occur because the optical measurements force the AFKF to move the mount away from the target.

To simulate an offset, or registration, error between the radar and optical sensor, a 5 milli-radian azimuth offset was added to the optical measurements. The plot in figure 10 shows the effect on the track azimuth, where the dashed line represents the radar Kalman filter track, the dotted the optical Kalman filter track, and the solid the AFKF track. The AFKF track falls between the two sets of measurements, its actual position determined by the relationship between the noise covariance of the two sensors. The AFKF track will be closer to the sensor with the lowest noise covariance; therefore if the optical sensor is correctly aligned, the error is relatively small. However if the radar is correctly aligned, and the optical sensor is misaligned, the error is much larger. Therefore it is important to align the higher resolution sensors accurately.

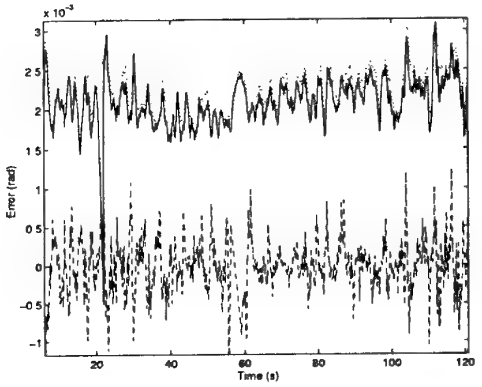

Figure 10: Azimuth tracking error with optical misalignment.

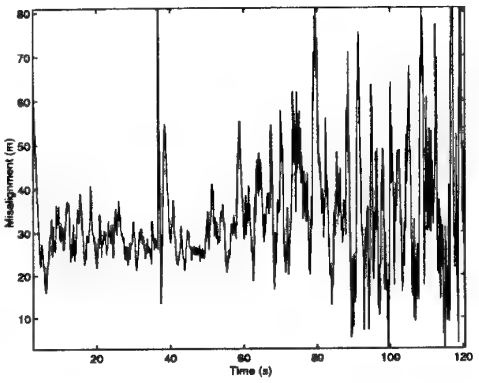

Figure 11: Track error caused by jet vapour trail.

The test-bed attempted to track a Boeing 737 commercial aircraft producing a distinct vapour trail. The radar maintained track on the body of the aircraft. However the optical tracker could not maintain a steady centroid position on the target, preferring to move randomly along the aircraft body and the vapour trail. As the optical measurements dominate the AFKF, the fused track also moved along the vapour trail, creating large track errors. Figure 11 shows the error between the AFKF and the radar Kalman filter tracks. The fused track lags an average of 30 metres behind the radar, reaching distances of up to 80 metres at times. In particular, the last half of the track shows large excursions along the vapour trail.

Two aircraft were observed flying elliptical circuits in formation. While tracking, the radar frequently switched between targets. This is particularly likely when targets are separated in azimuth but not range, because the radar's angular discrimination is relatively poor. The higher angular resolution of the optical system enabled it to maintain track on a single aircraft, but it often changed targets as they crossed in the field of view. Although this caused all three trackers to produce noisy tracks, they all maintained track on the formation. Therefore although the optical tracker may be able to discriminate between individual targets in a formation, it may not be able to maintain track on a particular target if the formation is manoeuvring.

The adaptive algorithm (section 5.3) was also evaluated using the real data. In general its performance was similar to that of the AFKF. However it showed a lower tendency to be seduced by interfering targets. Under certain conditions, the predicted target position does not immediately follow the interfering target. This causes the estimate of optical measurement noise covariance to increase and reduce the contribution of the optical sensor. This allows the AFKF to follow the radar measurements and continue tracking the target. However the number of lost tracks due to seduction of the optical sensor was still unacceptable.

8 Model Mismatch

8.1 Overview

Successful tracking requires a number of assumptions, particularly with regard to the expected dynamics of the target and the characteristics of the measuring sensors. This knowledge is expressed using dynamic and measurement models respectively. Of particular interest to the tracking filter designer is the effect of mismatch between the assumed models and the actual scenario. This is often referred to as sensitivity analysis. Sensitivity analysis has been covered extensively for single sensor systems, and in particular for the state space based Kalman filter [26].

This section evaluates the effect of adding an additional sensor on the robustness of the tracking filter to model mismatch, or alternatively, on the sensitivity of the tracking filter to model errors.

8.2 Sources of Model Mismatch

Model mismatch may occur in either the state transition matrix or the measurement matrix [26]. For linear-Gaussian filtering, the dynamic model mismatch may occur in the state transition matrix (F). However, any mismatch here is generally approximated by additional process noise. The assumption of Gaussian noise preserves the linearity of the filter, and the effect of non-Gaussian noise is not considered. Also the effect of bias on this noise is ignored. The process noise covariance is a critical design parameter of the Kalman filter, and the effect of errors in its value is considered. If the assumed value is too high, the filtered state estimate will contain excessive noise. A value which is too low will ultimately cause divergence from the true track.

The measurement function (H) is another potential source of model mismatch. This model may be a linear approximation of a non-linear coordinate transformation. Again this may be approximated by additional measurement noise. As for the process noise, the effect of non-Gaussian noise is not considered, and it is assumed that any bias can be effectively removed through correct sensor registration. The measurement noise covariance is also a critical design parameter. An excessive design value will reduce the Kalman gain, increasing the possibility of track divergence. A smaller value will increase the gain, introducing additional sensor noise into the state estimate.

8.3 Evaluation

8.3.1 Performance measures

To evaluate the effect of the measurements from the additional sensor, it is first necessary to remove the effect of fusion from the track error covariance. Assuming that the average or mean tracking error is zero, a useful measure of tracking performance is the covariance of the tracker's position errors, or *track error covariance.* [1] The lower the track error covariance, the better the tracking performance. The improvement in tracking due to sensor fusion can then be defined as the ratio of the track error covariance of a single sensor tracker to the track error covariance of the dual sensor tracker, referred to here as the *fusion gain* (FG). If the FG is determined when both the single and dual sensor trackers are operating under matched conditions, it is referred to as the *matched* FG. A FG greater than unity implies an improvement in tracking performance, and a value below unity, a degradation of tracking performance.

The *mismatched* FG provides a measure of the change in tracking performance of a mismatched tracker caused by the addition of another sensor. Of particular interest is the comparison between this and the corresponding change under matched conditions. A *mismatch factor* (MF) is introduced as the ratio of mismatched FG to matched FG. The MF is a measure of the improvement obtained in the tracking performance of a mismatched tracker by adding a second sensor, ignoring the expected improvement due to the fusion process alone. This is effectively a measure of the improvement in the tracker's sensitivity to model mismatch, with values greater than unity indicating reduced sensitivity to model mismatch, and values below unity indicating increased sensitivity.

8.4 Results

8.4.1 Simulations

A selection of single sensor and dual sensor Kalman filters with various design values of process noise covariance and measurement noise covariance were applied to two hundred sets of statistically equivalent data. The average steady state track error covariance was determined for each filter.

The effect of process model mismatch was evaluated by varying the design value of the filter process noise covariance while matching the measurement noise covariance to the data. Similarly, the effect of measurement noise mismatch was evaluated by varying the filter design value of the measurement noises while maintaining matched process noise.

8.4.2 Process Noise Mismatch

The results in figure 12 show that the MF is approximately unity for design process noise covariances over the range from 0.001 to 1000 times that of the data. This indicates

[1] Under zero average track error conditions, the track error covariance is defined as the average squared track error.

that the contribution of the additional sensor to the sensitivity of the system to process noise covariance mismatch is insignificant.

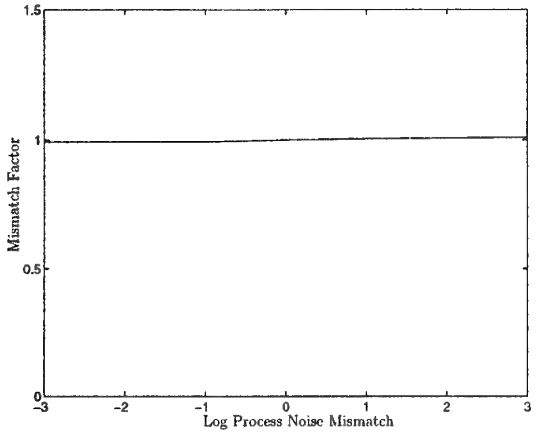

Figure 12: Mismatch factor for process noise mismatch

8.4.3 Measurement Noise Mismatch

This section considers various levels of model mismatch in both sensors. The sensor used in the single sensor system is referred to as the *original sensor*, and the sensor added to it to form the dual sensor system is referred to as the *additional sensor*. The results are displayed as three dimensional meshes with an underlying contour plot for clarity. The vertical axis indicates the value of either the logarithm of the matched FG, or the MF. The other two axes indicate the level of model mismatch in each sensor, where the scale represents the logarithm of the ratio of the filter design measurement noise covariance to the actual measurement noise covariance of the data, eg 2 denotes a design covariance of 100 times that of the data, and zero denotes the matched condition. A detailed discussion of these results may be found in [29].

The mismatch factor for equal measurement noise covariance in both sensors is shown in figure 13. These results show that fusing the measurements from an additional sensor whose design measurement noise covariance is greater than or equal to the original sensor (ie towards the right of the plot), does not alter the system's robustness to measurement model mismatch. However, if the design noise covariance of the additional sensor is less than that of the original (ie towards the left of the plot), the fused system's performance is poorer than that expected from the filter design equations. Under these conditions the system becomes more sensitive to modelling errors.

Figure 14 illustrates the results when the data received from the additional sensor has a noise covariance ten times that of the original. It can be seen that these results are similar to those obtained when the measurement noise covariances are equal.

The results obtained when the noise covariance of the data from the additional sensor is ten times smaller than that of the original sensor are shown in figure 15. The MF is again similar to those obtained from the equal measurement noise covariances, with the

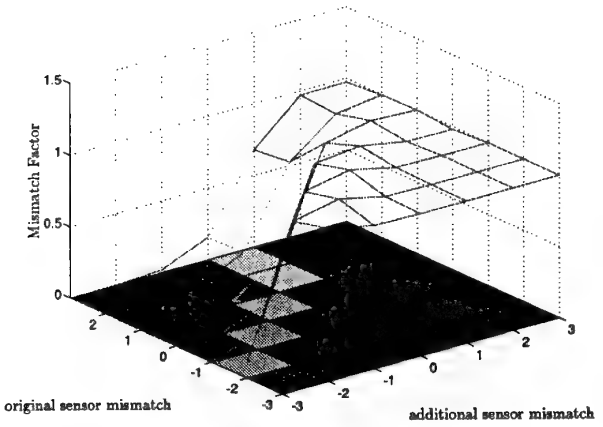

Figure 13: Mismatch factor for sensors with equal noise covariance.

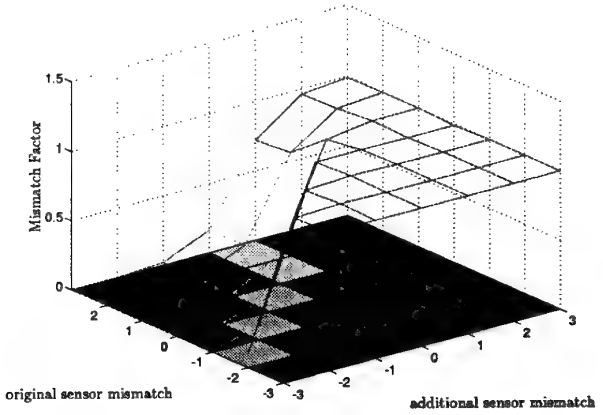

Figure 14: Mismatch factor when the additional sensor is noisier than the original.

exception of a peak of approximately 1.7 when the design measurement covariances are equal. These above unity values of MF indicate that, under these conditions, the system's sensitivity to model errors decreases.

9 Conclusions

An asynchronous fusion algorithm, the asynchronous fused Kalman filter (AFKF), has been developed using a variable update rate Kalman filter with an augmented measurement model. This algorithm has been evaluated using both simulated and real data. An adaptive version has also been evaluated.

A sensor fusion test-bed has been developed for multi-sensor data collection and al-

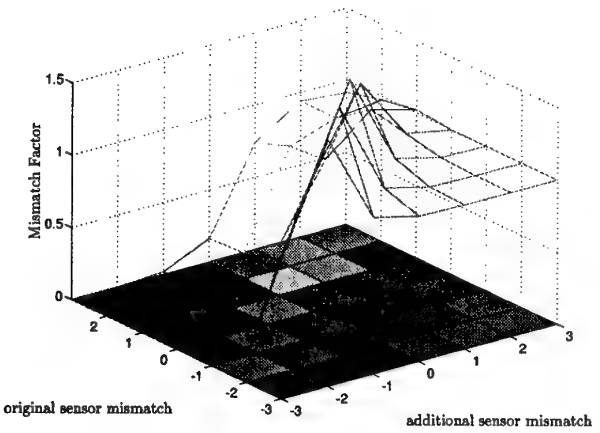

Figure 15: Mismatch Factor when the original sensor is noisier than the additional.

gorithm evaluation. It currently consists of two sensors, a pulse Doppler radar and an optical video tracker.

The asynchronous fused Kalman filter (AFKF) performed well under clear sky conditions. Its performance deteriorated at close range because of increased optical noise.

The AFKF performed well with intermittent sensor operation. When the measurements from one sensor ceased, the AFKF reverted to single sensor tracking, using the measurements from the remaining sensor.

The AFKF performed poorly in clutter and in the presence of interfering targets, often losing track. The optical sensor's susceptibility to clutter and other interfering targets was caused by its lack of range discrimination. To overcome this problem, some form of data association is required to determine which measurements should be used for tracking (eg Multi-Sensor Probabilistic Multi-Hypothesis Tracking [17, 18]).

When tracking in clutter or at close range, the ability to change the tracking filter models to suit the conditions, ie adaptive or multiple model techniques, would be advantageous.

The AFKF successfully tracked a formation of two aircraft. However, although the optical sensor was able to discriminate between the two aircraft, it frequently switched targets during manoeuvres.

The addition of another sensor does not affect the sensitivity of the tracking filter to errors in the noise covariance in the process or dynamic models. In general, the addition of another sensor only affects the sensitivity of the tracking filter to errors in measurement noise covariance when the tracking filter measurement noise covariance for the additional sensor is less than that for the original sensor. In this case, the system becomes more sensitive to errors in the measurement models.

References

1. Y Bar-Shalom and T E Fortmann, *Tracking and Data Association.* Academic, Orlando, Florida, 1988.

2. Y Bar-Shalom, editor, *Multitarget-Multisensor Tracking: Advanced Applications.* Artech House, Norwood, Massachusetts, 1990.

3. Y Bar-Shalom, editor, *Multitarget-Multisensor Tracking: Applications and Advances.* Artech House, Norwood, Massachusetts, 1992.

4. W D Blair, T R Rice, A T Alouani, and P Xia, "Asynchronous data fusion for target tracking with a multi-tasking radar and optical sensor," *Proc SPIE Acquisition, Tracking and Pointing V*, volume 1482, pages 234–245, Orlando, Florida, 1991.

5. A T Alouani and T R Rice, "On asynchronous data fusion," *Proc of 26th Southeastern Symposium on System Theory*, pages 143–6, Cookeville, Tennessee, Mar 1994.

6. Daniel W McMichael and Nickens N Okello, "Maximum likelihood registration of dissimilar sensors," *Proc 1st Australian Data Fusion Symposium*, pages 31–4, Adelaide, Australia, Nov 1996.

7. A T Alouani, T R Rice, and N Auger, "On the generalized nearest neighbour filter," *Proc of 22th Southeastern Symposium on System Theory*, pages 261–4, Cookeville, Tennessee, Mar 1990.

8. A Houles and Y Bar-Shalom, "Multisensor tracking of a maneuvering target in clutter," *IEEE Transactions on Aerospace and Electronic Systems*, AES-25(2):176–89, Mar 1989.

9. G W Pulford and R J Evans, "Probabilistic data association for systems with multiple simultaneous measurements," *Automatica*, 32(9):1311–6, Sep 1996.

10. D Musicki and R J Evans, "Tracking in clutter using probabilistic data association," *Proc of IEE International Conference Radar 92*, pages 82–5, Brighton, England, Oct 1992.

11. S Deb, R Mallubhatla, K R Pattipati, and Y Bar-Shalom, "A multisensor multitarget data association algorithm," *Proc of IEEE International Conference on Systems Engineering*, pages 320–3, Pittsburgh, Pennsylvania, Aug 1990.

12. S Deb, K R Pattipati, and Y Bar-Shalom, "A multisensor multitarget data association algorithm for heterogeneous sensors," *IEEE Trans on Aerospace and Electronic Systems*, AES-29(2):560–8, 1993.

13. S S Blackman and I P Bottlik, "Coordinated presentation of multiple hypotheses in multitarget tracking," *Proc SPIE Signal and Data Processing of Small Targets*, volume 1096, pages 152–9, Orlando, Florida, Mar 1989.

14. S S Blackman, R J Dempster, and J T Fagarasan, "Continuous time representation of multiple hypothesis track data," *Proc SPIE Signal and Data Processing of Small Targets*, volume 1954, pages 352–60, Orlando, Florida, Apr 1993.

15. D Avitzour, "A maximum likelihood approach to data association," *IEEE Trans on Aerospace and Electronic Systems*, AES-28(2):560–5, Apr 1992.

16. R L Streit and T E Luginbuhl, "Maximum likelihood method for probabilistic multi-hypothesis tracking," *Proc SPIE Signal and Data Processing of Small Targets*, volume 2235, pages 394–405, Orlando, Florida, apr 1994.

17. M L Krieg and D A Gray, "Track fusion in the presence of an interference," *Proc 4th International Symposium on Signal Processing and its Applications*, volume 1, pages 192–5, Gold Coast, Australia, Aug 1996.

18. M L Krieg and D A Gray, "Multi-sensor probabilistic multi-hypothesis tracking," *Proc 1st Australian Data Fusion Symposium*, pages 153–8, Adelaide, Australia, Nov 1996.

19. Evangelos Giannopoulos, Roy Streit, and Peter Swaszek, "Probabilistic multi-hypothesis tracking in a multi-sensor, multi-target environment," *Proc 1st Australian Data Fusion Symposium*, pages 184–9, Adelaide, Australia, Nov 1996.

20. L Hong, "Adaptive data fusion," *IEEE International Conference on Systems, Man, and Cybernetics - Decision Aiding for Complex Systems*, volume 2, pages 767–72, Charlottesville, Virginia, Oct 1991.

21. L Hong, "Adaptive multi-sensor integration in uncertain environment with imperfect sensors," *Proc 30th IEEE Conference on Decision and Control*, volume 1, pages 142–3, Brighton, England, Dec 1991.

22. X Xie and R J Evans, "Multiple target tracking using hidden markov models," *Proc of IEEE 1990 International Radar Conference*, pages 625–628, Arlington, Virginia, May 1990.

23. F Martinerie and P Forster, "Data association and tracking from distributed sensors using hidden markov models and evidential reasoning," *Proc 31st IEEE Conference on Decision and Control*, volume 4, pages 3803–4, Tuscan, Arizona, Dec 1992.

24. F Martinerie and P Forster, "Data association and tracking using hidden markov models and dynamic programming," *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, volume 2, pages 449–52, San Francisco, California, Mar 1992.

25. A H Jazwinski, *Stochastic Processes and Filtering Theory.* Academic, San Diego California, 1970.

26. A Gelb, *Applied Optimal Estimation.* MIT Press, Cambridge, Massachusetts, 1992.

27. R K Mehra, "On the identification of variances and adaptive kalman filtering," *IEEE Trans on Automatic Control*, AC-15:175–84, Apr 1970.

28. M L Krieg and D A Gray, "Radar and optical track fusion using real data," *Proc 1st Australian Data Fusion Symposium*, pages 25–30, Adelaide, Australia, Nov 1996.

29. M L Krieg and D A Gray, "Performance of state space multi-sensor track fusion with model mismatch," *Proc 1st Australian Data Fusion Symposium*, pages 1–6, Adelaide, Australia, Nov 1996.

Appendix A
Generic Pulse Doppler Radar Specification

Antenna

Prime focus fed amplitude monopulse	
Diameter	1500mm
Gain	39 dB
3 dB sum beamwidth	1.6°

Antenna Mount

Analogue rate loops in Azimuth and Elevation	
Small signal rate rise time	10 ms (az)
	15 ms (el)
Angle resolution	0.383 mrad (14 bit)

Transmitter

Peak power	200 W
Max duty cycle	30%

Receivers

Three matched channels	
Operating frequency	9.0 – 10.0 GHz
Noise figure	< 5 dB
Analogue to digital converter resolution	12 bits
Max sample rate (I and Q data)	8×10^6 samples/s
AGC range	78 dB
AGC resolution	2 dB

Timing and Control

Pulse repetition interval (PRI)	5 – 255 μs
Pulse length	1 μs
Block length (coherent interval)	1.4 – 65 ms
Range gate resolution	62.5 ns

Tracking parameters

Target range cells	2 per channel
Noise range cells	4 in sum channel
Doppler filters per range cell	128
Pulse repetition frequency (PRF)	13.157 – 26.315 kHz
Filter cycle time	33.3 ms

Appendix B
Video Camera Specification

Configuration

Colour CCD video camera with 75 mm automatic iris lens and 2× extender

Field of view (FOV)	2.48° × 1.87°

Camera

Pick-up device	$\frac{1}{2}$" colour, interline-transfer CCD
Picture element	681 (H) × 582 (V)
Scanning area	6.5 mm (H) × 4.9 mm (V)
Scanning system	2:1 interlaced
Scanning frequency	15.625 kHz (H) 50 Hz (V)
Resolution	> 430 TV lines (H) > 420 TV lines (V)
Signal to noise ratio (SNR)	> 48 dB
Output	VBS 1.0 Vp–p PAL compatible

Appendix C
Adept20 Automatic Video Tracker Specification

Video Input

Format	Composite Video 635/525 line
Level	1.0 Vp–p
Standards	CCIR or RS170
Type	Differential
Impedance	75 Ω

Video Output

Normal video with symbology overlay

Format	Composite Video
Level	1.0 Vp–p (into 75 Ω)
Impedance	75 Ω

Video Preprocessing

Spatial enhancement
Statistical enhancement
Threshold (positive, negative contrast)

Track Algorithms

Centroid
Correlation

Track Window Position

Manual
Automatic (follows target)

Track Window Size

Manual
Adaptive (automatically encloses target)

Breaklock/Coast

Automatic breaklock algorithm
Automatic target re-acquisition after breaklock

Performance

Typical for target size $> 6 \times 6$ TV lines	
Min target contrast	$< 5\%$
Min signal to noise ratio	< 4
Noise on target position output	< 1 TV line (3σ)
Error update rate	50 Hz
Error latency	< 15 ms

VMEbus Interface

Capability	Slave, A24, D08(O)
Board select	single input
	2 boards hosted on a backplane

Serial Interface

Electrical	RS232 or RS422
Format	asynchronous
Baud rate	9600

THIS PAGE IS INTENTIONALLY BLANK

DISTRIBUTION LIST

ASYNCHRONOUS SINGLE PLATFORM SENSOR FUSION

Mark L Krieg

	Number of Copies
DEFENCE ORGANISATION	
Task Sponsor	
Director General Force Development (Aerospace)	1
SYS1, DDSYS	1
S&T Program	
Chief Defence Scientist FAS Science Policy AS Science Corporate Management	1
Counsellor, Defence Science, London	Doc Data Sht
Counsellor, Defence Science, Washington	Doc Data Sht
Scientific Adviser to MRDC, Thailand	Doc Data Sht
Director General Scientific Advisers and Trials Scientific Adviser Policy and Command	1
Navy Scientific Adviser	Doc Data Sht
Scientific Adviser, Army	Doc Data Sht
Air Force Scientific Adviser	1
Director Trials	1
Aeronautical and Maritime Research Laboratory	
Director, Aeronautical and Maritime Research Laboratory	1
Electronics and Surveillance Research Laboratory	
Director, Electronics and Surveillance Research Laboratory	1
Chief, Microwave Radar Division	1
Research Leader, Microwave Radar	1
Head, Radar Processing	1
Task Manager	1
Author	1
DSTO Libraries	
Library Fishermens Bend	1
Library Maribyrnong	1
Library Salisbury	2
Australian Archives	1
Library, MOD, Pyrmont	Doc Data Sht

Forces Executive

Director General Force Development (Sea)	Doc Data Sht
Director General Force Development (Land)	Doc Data Sht

Navy

SO(Science), Director of Naval Warfare, Maritime Headquarters Annex, Garden Island	1

Army

ABCA Office, G-1-34, Russell Offices, Canberra	4
SO(Science), HQ 1 Division, Milpo, Enoggera, Qld 4057	1

Air Force

Chief Weapons Systems Instructor, School of Air Navigation	1

S&I Program

Defence Intelligence Organisation	1
Library, Defence Signals Directorate	Doc Data Sht

B&M Program(libraries)

Officer in Charge, TRS, Defence Central Library	1
Officer in Charge, Document Exchange Centre	1
Additional copies for DEC for exchange agreements	
US Defense Technical Information Center	2
UK Defence Research Information Centre	2
Canada Defence Scientific Information Service, Canada	1
NZ Defence Information Centre, New Zealand	1
National Library of Australia	1

UNIVERSITIES AND COLLEGES

Australian Defence Force Academy Library	1
Head of Aerospace and Mechanical Engineering, ADFA	1
Librarian, Flinders University	1

OTHER ORGANISATIONS

NASA (Canberra)	1
Australian Government Publishing Service	1
The State Library of South Australia	1
Parliamentary Library of South Australia	1

ABSTRACTING AND INFORMATION ORGANISATIONS

INSPEC: Acquisitions Section Institution of Electrical Engineers	1
Library, Chemical Abstracts Reference Service	1
Engineering Societies Library, US	1
Materials Information, Cambridge Science Abstracts	1
Documents Librarian, The Center for Research Libraries, US	1

INFORMATION EXCHANGE AGREEMENT PARTNERS

Acquisitions Unit, Science Reference and Information Service, UK	1
Library – Exchange Desk, National Institute of Standards and Technology, US	1
National Aerospace Library, Japan	1
National Aerospace Library, Netherlands	1

SPARES

DSTO Salisbury Research Library	10
Total number of copies:	61

DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION DOCUMENT CONTROL DATA	1. CAVEAT/PRIVACY MARKING

2. TITLE	3. SECURITY CLASSIFICATION
ASYNCHRONOUS SINGLE PLATFORM SENSOR FUSION	Document (U) Title (U) Abstract (U)
4. AUTHOR(S)	5. CORPORATE AUTHOR
Mark L Krieg	Electronics and Surveillance Research Laboratory PO Box 1500 Salisbury, South Australia, Australia 5108

6a. DSTO NUMBER	6b. AR NUMBER	6c. TYPE OF REPORT	7. DOCUMENT DATE
DSTO–TN–0084	AR–010–219	Technical Note	June, 1997

8. FILE NUMBER	9. TASK NUMBER	10. SPONSOR	11. No OF PAGES	12. No OF REFS
Z7208/3/100	ADA93/021	DGFD(AERO)	29	29

13. DOWNGRADING / DELIMITING INSTRUCTIONS	14. RELEASE AUTHORITY
Not Applicable	Chief, Microwave Radar Division

15. SECONDARY RELEASE STATEMENT OF THIS DOCUMENT

Approved For Public Release

OVERSEAS ENQUIRIES OUTSIDE STATED LIMITATIONS SHOULD BE REFERRED THROUGH DOCUMENT EXCHANGE CENTRE, DIS NETWORK OFFICE, DEPT OF DEFENCE, CAMPBELL PARK OFFICES, CANBERRA, ACT 2600

16. DELIBERATE ANNOUNCEMENT

No Limitations

17. CITATION IN OTHER DOCUMENTS

No Limitations

18. DEFTEST DESCRIPTORS

Multisensors
Sensor data fusion
Kalman filtering
Performance evaluation

19. ABSTRACT

Multi-sensor tracking potentially has many advantages over single sensor tracking. This report evaluates the performance of a multi-sensor tracking algorithm, the asynchronous fused Kalman filter, using both simulated and real data from two dissimilar sensors. The real data was collected using a sensor fusion test-bed consisting of two sensors, a pulse Doppler radar and an optical video tracker. The performance of the algorithm has been evaluated under various conditions including clear sky, clutter, multiple targets and intermittent sensor operation. The effect of sensor fusion on the system's robustness to model mismatch has also been investigated.